

Final Performance Report  
Concordia Final Report  
Grant number: PX-50003-08  
Project Title: Concordia  
Project Director: Dr. Roger Bagnall  
Grantee Institution: New York University  
Date Submitted: 30 June 2010

## 1 Introduction: Aims, Scope and Audience

The Concordia project was a joint effort of the Institute for the Study of the Ancient World (ISAW) at New York University and of King's College London. It was designed to explore and demonstrate ways to link together scholarly resources published separately on humanities reference websites. For trial purposes, we identified a small number of discrete, digital collections in development or planning stages at the partner institutions.<sup>1</sup> These collections hold significant value for the study of antiquity and, by virtue of their geographic coverage and heterogeneous nature, reflect the project team's focus on presentation of variegated content drawn from multiple sources in ways that straddle conventional, entrenched scholarly boundaries.

Some of these collections comprise reference resources, providing information about ancient settlements and geographic features. Others address ancient text-bearing objects (inscriptions and papyri), variously including descriptions, images, translations, commentary and even the edited texts themselves. Geographically, they span areas normally studied by different specialists. Their combined footprint constitutes a contiguous but varied stretch of the ancient world, covering more than 830,000 square kilometers: from Tripolitania (southern Tunisia and western Libya) along the coast eastward to the Nile delta and then southward up the river to include the Fayum and several of the best documented ancient administrative regions (the Oxyrhynchite, Kynopolite, Herakleopolite and northern Arsinoite nomes). The work plan in our joint proposals to NEH and JISC set aside significant time for fleshing out the collections themselves to ensure the necessary geographic and contextual overlap.

In seeking to identify and digitally express the relationships between content in these collections, our goal was to move beyond the ubiquitous, plain hypertext links that make it possible for human users to click and jump easily from one web site to the next. We wanted *typed* links: connections between resources that were labeled to indicate the reason the link was being made (e.g., to indicate the findspot of an inscribed document) so that automated agents could extract the link types for subsequent analysis. We planned to develop or select a succinct vocabulary of those relationship types and to experiment with standards-based methods for encoding the relationships. The process whereby we arrived

---

<sup>1</sup> Collections selected for Concordia work included: the *Duke Databank of Documentary Papyri* (DDBDP: <http://idp.atlantides.org/trac/idp/wiki/DDBDP>), the *Heidelberger Gesamtverzeichnis der griechischen Papyrusurkunden Ägyptens* (HGV: <http://idp.atlantides.org/trac/idp/wiki/HGV>), the *Advanced Papyrological Information System* (APIS: <http://idp.atlantides.org/trac/idp/wiki/APIS>), the *Pleiades Project* (<http://pleiades.stoa.org>), the *Inscriptions of Aphrodisias 2007* (IAph2007: <http://insaph.kcl.ac.uk/iaph2007/index.html>), and the *Inscriptions of Roman Tripolitania 2009* (IRT2009: <http://irt.kcl.ac.uk/irt2009/>).

at the terms in the “Concordia Thesaurus” and at the selection of the Atom Syndication Format for encoding are described below in section 2.1.<sup>2</sup>

All of the textual collections involved in the project employ the “EpiDoc” standard for the encoding of their content.<sup>3</sup> EpiDoc (a name drawn from the standard’s origin as a method for encoding epigraphic documents) is a fully valid customization (subset) of the Text Encoding Initiative (TEI) Guidelines for Electronic Text Encoding and Interchange (version P5).<sup>4</sup> Wikipedia describes the TEI fairly thus:

Since 1994, these guidelines are a widely-used standard for text materials for performing online research and teaching, and TEI is now the de facto standard for the encoding of electronic texts in the humanities academic community.<sup>5</sup>

The Concordia approach to the use of feeds and terminological thesaurus does not presuppose or require collection content to be encoded in EpiDoc (indeed, Pleiades makes no use of EpiDoc for its geographic information); however, the fact that all of the textual materials used for this project were encoded in EpiDoc facilitated the reuse of tooling designed to produce the feeds (principally eXtensible Stylesheet Language Transformations, XSLTs).<sup>6</sup>

The target audience of the project was varied and international. Scholars planning or conducting projects that will produce online digital resources may find the Concordia model useful. Both students and scholars in the United States, Europe and North Africa have much to gain from the digital availability of the content we have assembled and linked. The project study area – whether understood geographically, temporally or by content type – spans a number of subdisciplines and national communities. Seamless access, a goal toward which Concordia moves us, promises to facilitate better dialog and information exchange between these constituencies.

---

<sup>2</sup> Concordia Thesaurus:

<http://www.atlantides.org/trac/concordia/wiki/ConcordiaThesaurus>; Atom: M. Nottingham and R. Sayre (eds.), *The Atom Syndication Format*, Network Working Group Request for Comment (RFC) 4287, December 2005, <http://tools.ietf.org/html/rfc4287>.

<sup>3</sup> EpiDoc: <http://epidoc.sf.net>

<sup>4</sup> TEI: <http://www.tei-c.org>

<sup>5</sup> *Wikipedia: The Free Encyclopedia*, s.v. “Text Encoding Initiative” by ‘Schnolle’ et al., version of 26 June 2010 at 03:27, [http://en.wikipedia.org/wiki/Text\\_Encoding\\_Initiative](http://en.wikipedia.org/wiki/Text_Encoding_Initiative)

<sup>6</sup> XSLT: James Clark (ed.), *XSL Transformations (XSLT), Version 1.0*, W3C Recommendation 16 November 1999, <http://www.w3.org/TR/xslt>; Michael Kay (ed.), *XSL Transformations (XSLT), Version 2.0*, W3C Recommendation 23 January 2007, <http://www.w3.org/TR/xslt20/>.

## 2 Project Activities and Accomplishments

The original proposal divided tasks into five discrete areas, each of which is addressed below. During the course of the project, we also identified a number of additional tasks that proved essential to the conduct of the effort. These are described in section 2.6.

### 2.1 Create "web feeds for interoperability" for existing and new digital publications

The Concordia team set out to define a basis for interoperability between the digital publications involved in the project. The goal was to establish Atom web feeds for the content in each publication. These feeds were to include links to resources in other digital publications, and the links were to be annotated using a terminological thesaurus in order to indicate relationships of scholarly interest, such as "place of finding" or "evidence for a place-name". Tooling (in the form of XSL style sheets) was promised for the conversion of content to Atom, thereby retrofitting these publications to support interoperable processes that consume web feeds and provide discovery and linking services on the basis of information found therein.

Our original idea that feeds could act as a transport medium for summaries of and linkages between published articles and objects has been borne out by our experiences. Likewise, the idea that we needed our own link-type thesaurus also proved true. We have also concluded that these decisions need to be made across a large community. The point of our work was to surface information that is already in effect in the source documents, but to do it in a way that would have potentially large application across a variety of projects.

We also clarified the definition and goals of Concordia with respect to "federated search" applications. Concordia approaches search from the web search engine angle. We are building an index from the data collected by a web crawler, and not creating a federated search space that is invisible to other search engines on the web.

Our original objectives included a phased release of web feeds for *Inscriptions of Aphrodisias*, *Inscriptions of Roman Tripolitania*, the *Inscriptions of Roman Cyrenaica* and the Integrating Digital Papyrology (IDP) community (DDB, HGV and APIS). The *I Aph* feeds were generated in the first instance, based on online, already published material. Due to the refocusing of work on the *IRT* texts and infrastructure to enable re-use by the *IRCyr* project, the *IRCyr* texts are not yet available in XML form and so the generation of web feeds for these texts has been deferred beyond the end of this project. The XSLT that was used to generate the *I Aph* and *IRT* feeds will, however, be largely suitable for this final task as it stands. Because of the delays incurred in transfer of responsibility for the Papyrological Navigator (PN) software and for its refactoring, production of feeds for the IDP materials has also been deferred; however, example feeds were prepared and vetted.

See section 4, "Grant Products" for a link to web addresses of the various feeds.

#### 2.1.1 The Concordia Thesaurus

Significant project time was devoted to discussions of the thesaurus and its content. These discussions included members of our advisory board, who themselves represent a wide

array of projects and publications that contain resources that could eventually be linked with the Concordia collections. Ultimately, we elected to produce our own thesaurus, incorporating only those terms necessary for immediate needs of the project. We rejected the idea of trying to produce a “complete” or “universal” thesaurus for linking ancient studies resources – though the matter was discussed at length during our second workshop – on the grounds that a much wider process of consultation and consensus-building would be needed than the present grant could support.

Our decision to produce our own thesaurus was informed by analysis of the CIDOC Conceptual Reference Model (CRM), and in particular its extensive vocabulary of relationships used in cultural heritage documentation.<sup>7</sup> This vocabulary is in use in a number of museum contexts, and is also being employed in the latest redesign of the Lexicon of Greek Personal Names (LGPN).<sup>8</sup> At the end of an extensive discussion and joint modeling exercise during our second workshop, Concordia team members concluded that there were a number of challenges inherent in trying to incorporate CIDOC CRM concepts and terms into our relationship model. Further consideration of these challenges were discussed and explored after the workshop, culminating in a decision in December 2008 that CIDOC CRM relationship models and terms would not be adopted for Concordia at this time. The rationale underlying this decision, as well as a concise explanation of Concordia’s goals and methods, were published in an online white paper.<sup>9</sup>

### **2.1.2 Atom Syndication Format**

The original proposal identified the Atom feed format as the best carrier for metadata concerning the content of individual resources and the relationships between them and resources in other projects. This assumption proved valid, but we arrived at that conclusion only after evaluating two alternative formats.

The first alternative format, the Open Archives Initiative Object Reuse and Exchange (OAI/ORE) specification, was originally published concurrently with the start of our project.<sup>10</sup> It defines standards for the description and exchange of aggregations of Web resources. Although it initially appeared to us that OAI/ORE would be useful – not least because its default implementation was defined as Atom feeds (using special conventions) – we ultimately decided that we could express all of the relationships we needed to via plain Atom feeds, without the added OAI/ORE conventions. Our rationale was two-fold. Firstly, OAI/ORE incorporates a number of “beyond Atom” features that are good for institutional repositories but that confer no extra value for interoperation between the Kings digital epigraphic publications and Pleiades (e.g., packaging and aggregation of part-of works). Secondly, the major search engines index Atom but have no special indexing capacity for OAI/ORE; this would push the relationships between our resources into the hidden web, and it is one of our fundamental tenets that what we do should be broadly discoverable on the open web.

---

<sup>7</sup> CIDOC-CRM: <http://cidoc.ics.forth.gr/>

<sup>8</sup> Lexicon of Greek Personal Names: <http://www.lgpn.ox.ac.uk/>

<sup>9</sup> See Section 4, “Grant Products”, sub “White paper: Gillies, Sean”.

<sup>10</sup> OAI/ORE: <http://www.openarchives.org/ore/>

We also addressed the relevance of the Resource Description Framework (RDF) and web-crawling for "social" networks of object/name/place data.<sup>11</sup> We concluded that because our new vocabulary of link relations bears some resemblance to a group of RDF predicates, Concordia graphs expressed in Atom could be transformed to RDF serializations for Semantic Web functionality in future. As far as we know, only one major ancient studies project in our immediate orbit is using RDF/Semantic Web for this purpose (nomisma.org).<sup>12</sup>In our opinion, there is not yet a widely endorsed use case for RDF in ancient studies, but we can reserialize our feeds in future if it becomes highly valuable. This is one of many reasons that we shall continue to discuss these issues in blogs and other online fora. The project site concordia.atlantides.org will continue to serve as a platform for publication and ongoing discussion of work in this area (for instance the DFG/NEH-funded *Epigraphic Interoperability Workshops*).

## 2.2 Establish web services to exploit feeds

The Concordia team set as a key goal the creation of web services to make use of the planned web feeds. The primary objective was discovery services, exploiting two aspects of the anticipated connections between resources. Firstly, we intended to develop a limited-domain web crawler that would regularly harvest new feed content from the participating publications and use it to build an index. This index would be used by a search service designed to understand the scholarly distinctions—like "findspot" or "observed location"—embedded in the feed links. Secondly, we aimed to exploit geographic information stored in Pleiades to facilitate geographic search and visualization for all publications linked to it using our web feed implementation. Although the realization of an in-production service family proved too ambitious for the time frame and funding level of the Concordia project itself, we did produce a prototype toolset that will be built on and employed across ISAW's digital projects portfolio in future.

The accompanying figure (below) illustrates the major components and data flows of the experimental web crawler and relationship search engine designed under Concordia funding. We have continued development with in-house funds during the last year. The crawler (yellow) maintains a search "horizon", essentially a list of web resources to index.<sup>13</sup> When the crawler cycles, it fetches one of these resources, parses Concordia-compatible typed links out of it, puts those links into the horizon collection and then indexes the content of the resource. The horizon list is seeded initially with the known locations of the feeds discussed above. Additional start points can be added manually in future, or automatically discovered via Concordia-style links in resources previously identified. For each resource fetched by the crawler, indexed information is written to disk (blue) in Atom format, ready for supervised loading into our search database (green). This object-oriented

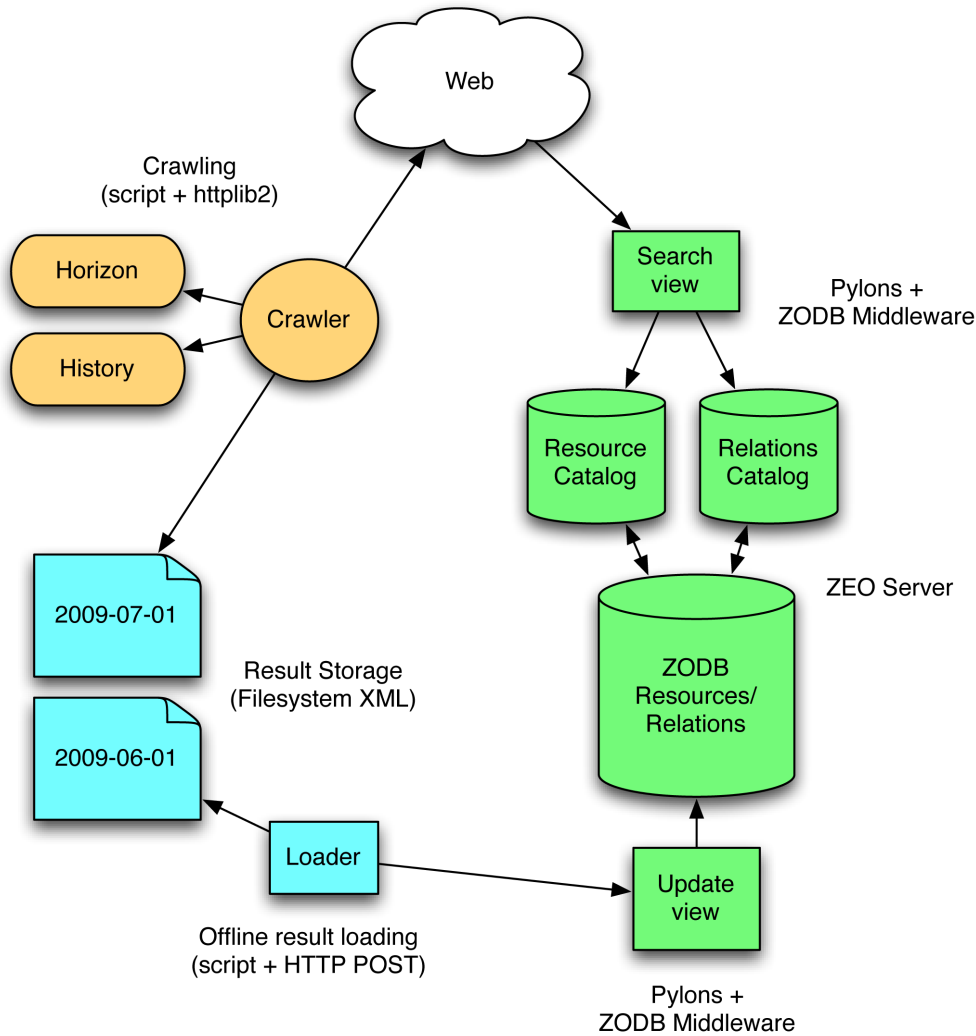
---

<sup>11</sup> Resource Description Framework: <http://www.w3.org/RDF/>

<sup>12</sup> <http://nomisma.org/>

<sup>13</sup> Our crawler design follows proven patterns laid out in: Pant, G., Srinivasan, P. & Menczer, F., 2004. "Crawling the Web." In M. Levene & A. Poulouvasilis, eds. *Web dynamics: adapting to change in content, size, topology and use*. Berlin: Springer, pp. 153-178. Available at: <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.14.1830>.

database provides search catalogs for swiftly querying key attributes of the resources we have indexed and the typed relationships they assert among resources. A search view provides mechanisms for querying this database and returning links to the resources matching those queries.



The Concordia work plan envisaged 25% of the ISAW technical work being spent on Pleiades upgrades and site maintenance. Indeed, during the period of performance, Pleiades saw a number of software improvements aimed at supporting better user interaction and at fitting its data model to the needs of the other Concordia team members. These improvements were rolled out during summer 2009. One major advance is the implementation of "part of" relationships between places: a more general and flexible development from the "feature" records capability outlined in the spring 2009 interim report. This enhancement permits the identification and mapping (where coordinates are known) of monuments, structures and areas within sites in order to provide more specific contexts for the indication of findspot, original location and the like. All places and names

are uniquely identified by stable URIs, thus providing the necessary references for feed-based linkages.

We have also gone through a major software migration, upgrading our platform from Plone 2.5 to Plone 3, which required revisions to most of our custom plug-in software and a replacement of our entire deployment apparatus. The extent of work required for the version upgrade was unanticipated, a consequence of major architecture and tooling changes introduced by the Plone software development team. This transition, however, gave us the opportunity to significantly improve our map interface and to implement a server-side caching strategy that has resulted in faster performance.

### **2.3 Reconfiguration of the Papyrological Navigator for search and display of arbitrary EpiDoc content**

The Concordia team also aimed to test the fitness of a separately developed papyrological web application for the search and display of epigraphic information, such as that contributed to the project by the King's team. The "Papyrological Navigator" (PN) was prototyped by Columbia University and was slated for release under the terms of the GNU Public License in July 2008. Our objective was to install a copy of the software on our servers at NYU, attempt to load the Concordia epigraphic materials into it, evaluate the functioning of its substring and lexical search capabilities and assess how it might be augmented to work with the feed-based services being developed by Concordia.

The PN software was not publicly released in July 2008 as expected. Delays—external to the Concordia Project Team—arose from Columbia's decision to out-source work on the Papyrological Navigator. NYU's Digital Library team ultimately emerged as the subcontractor for this work, but was not able to hire an appropriately skilled programmer until February 2009. We therefore deferred Concordia work on the PN until this programmer was in place and the software had been transferred from Columbia to NYU (summer 2009). For this reason, we requested (and received) a no-cost extension to the NEH grant in order to accommodate this work.

Once the PN was deployed at NYU, the Concordia team's analysis of the software's suitability for reuse with epigraphic materials could begin. It soon became clear that the PN would require significant redesign to support generalization of this type: an unexpected result. We offer the following analysis as a first step toward the necessary work.

The current version of the (PN) relies on a fairly heavyweight portal framework and a customized Apache Lucene search engine.<sup>14</sup> It is unable to simultaneously search documentary texts and metadata, and frequently has difficulty aligning texts in its constituent collections. A good deal of time and effort was expended on porting it over from Columbia to NYU, and although this was accomplished in the late summer of 2009, the project team decided that the system was too resource-intensive and ungainly. Moreover,

---

<sup>14</sup> The Papyrological Navigator is deployed at <http://papyri.info>. Lucene: <http://lucene.apache.org/java/docs/index.html>



work that was done on moving the APIS data over to NYU and converting it to EpiDoc in the Fall and Winter of 2009 meant that it would now be possible to better integrate the collections.

Accordingly, the PN is undergoing a redesign that will preserve its current functionality and also improve upon its capabilities. Because the redesign is still ongoing (though nearing completion), the method described here has not been thoroughly tested on a broad range of EpiDoc texts "in the wild." But inasmuch as it relies upon simple modifications to the standard EpiDoc XSLT stylesheets published on the EpiDoc Sourceforge site, it should work well with any document which conforms to the current schema.<sup>15</sup>

The PN search currently supports:

- Substring searches on source texts, including search on word fragments in the middle of a word. Word boundaries may indicated using an "anchor" (the "^" character). Search terms may be written in Beta Code or Unicode, and may account for or ignore both capitalization and diacriticals.
- Text searches may be filtered by place name or collection/publication.
- Lemmatized search of texts (i.e. search by dictionary headword).
- Metadata search, filtered by collection/publication, date, or place.

The new PN search supports these features and will in addition provide:

- Unified text, translation, and metadata search.
- Phrase and wildcard search.

The new PN search facilities are being provided by Apache Solr, a web-service wrapper around Lucene, which is the engine behind the current search.<sup>16</sup> It is very important to note that, while the current PN search relies upon a range of custom updates to Lucene, the sole customization to the new engine is the configuration of its search schema, an XML document that specifies what fields are to be indexed, and how. The various search features are enabled by indexing versions of the text with and without diacriticals and capitalization. These are converted from the source documents using an XSLT that generates a Solr add document (an XML document conforming to Solr's schema). This work means that there are now methods for indexing and searching EpiDoc using absolutely standard Open Source tools.

The mechanism for aligning documents from the PN's constituent collections is somewhat more customized, but relies again upon standard Open Source software. The collections surface links to other collections in the form of identifiers, and these are harvested and loaded into an RDF triplestore (Mulgara).<sup>17</sup> This triplestore can then be queried to locate related documents for any document in any of the collections. This is the mechanism used to drive the website generation and indexing for the new PN.

---

<sup>15</sup> EpiDoc stylesheets: <http://epidoc.svn.sourceforge.net/viewvc/epidoc/trunk/example-p5-xslt/>; EpiDoc schema: <http://www.stoa.org/epidoc/schema/8/tei-epidoc.rng>

<sup>16</sup> Apache Solr: <http://lucene.apache.org/solr/>

<sup>17</sup> Mulgara: <http://www.mulgara.org/>

Once these enhancements to the PN are deployed and validated by its user community (summer 2010), the software will be better fit for experiments such as those envisioned in the Concordia proposal. Both the PN interface and its search indexing components will no longer be so tailored for specific papyrological collections, facilitating more ready adaptation.

#### 2.4 4. Digitization, enhancement and dissemination of *The Inscriptions of Roman Tripolitania* (IRT)

This corpus of c. 1,000 texts was originally published by Joyce Reynolds and John Ward-Perkins in 1952 (British School at Rome); it has long been out of print, and was digitized (into HTML) a few years ago by a third party with the intention of distribution on CD ROM, but the project never came to publication. A particular aim was to add full illustration, which had not been possible in the circumstances of the early 1950s and remains prohibitively expensive in print today. We continued to work to this end, although the previously digitized images turned out to be of too low a standard: we therefore had to order a new set of high quality images (paid for externally). Our other principal enhancement was to link the publication to geographical data, which was not easily available to civilians in the 1950s. Our aim was to prepare this corpus for online publication during 2009. A further development was a decision to make the contents compatible with the Epigraphische Datenbank in Heidelberg, with a view to transferring data automatically to that widely used resource.<sup>18</sup>

The digital re-publication of *Inscriptions of Roman Tripolitania* (IRT) was carried out in several stages to maximize the efficiency and the extent to which the tools and processes can be repurposed for related projects in the future (including the ongoing *Inscriptions of Roman Cyrenaica* [IRCyr]). The existing digital copies of IRT, in HTML format, were acquired from the British School at Rome, and hand-edited by the King's team to remove any departures from the 1952 text of Reynolds and Ward-Perkins. The HTML was then converted to EpiDoc TEI XML using the CHET-C conversion tool, repurposing some of the code written for the Mellon-funded (2007-2008) *Integrating Digital Papyrology* project.

The texts were also brought somewhat up to date by the incorporation of texts, readings, and corrections both from the Addenda to the 1951 volume, and from a 1953 article. We deliberately limited our interventions to this level, in order to be able to complete the work within the time available. The TEI editions were then further enriched and enhanced by the addition of markup for indexing (words, names, places) and other metadata (categories, location history, dimensions) by a team of encoders including student interns. We found that this was of great benefit to the students as well as to us: they acquired a deeper understanding both of Roman inscriptions, and of the processes of digitization. All contributors were involved in studying the texts closely and also thinking about the metadata, and the various categories which encoding requires: this ensured that they engaged with the texts themselves, the historical content, and the processes.

---

<sup>18</sup> Epigraphische Datenbank Heidelberg: <http://www.uni-heidelberg.de/institute/sonst/adw/edh/>

Further work enriching the texts was carried out by means of script-assisted markup—in particular the lemmatization of Latin and Greek words using the Morpheus web service (very usefully on the servers of the Max Planck Institute in Berlin). The British School at Rome helped us to obtain high-quality scans of 1400 or so photographs to illustrate this edition.

The web process and design for the IRT publication was developed from the framework behind the *Inscriptions of Aphrodisias* and *IRCyr*. The improved tooling, indices and search functionality, and enhancements to the standard EpiDoc style sheets (including conversion scripts for creating database output for the Heidelberg Epigraphische Datenbank) were collaborative achievements with the Leverhulme-funded *IRCyr* project, and will also enhance the forthcoming *IRCyr* publication.

## 2.5 Digitization of five sets of map compilation materials and publication of the resulting data.

At the time the proposal was written, the Pleiades Community had brought the point features associated with two *Barrington Atlas* maps (including Cyrenaica) into the *Pleiades* format, and had enhanced the data with primary source references and original-script orthography (the original dataset and the *Barrington Atlas* employed a Roman character transliteration scheme for Greek names). Two additional maps were in work under separate funding. *Concordia* funds were budgeted to facilitate the digitization and incorporation of cultural features and names associated with five additional maps to provide seamless Gazetteer coverage for the project study area. The object of this expansion of *Pleiades* content was to provide geographic references for the epigraphic content to be produced in London, as well as the papyrological content expected to be released by the Integrating Digital Papyrology project. This work was completed on schedule. A full summary of current content (much has been added, under separate funding, since the *Concordia* materials) may be had on the Pleiades Content page.<sup>19</sup>

## 2.6 Unanticipated Tasks and Outcomes

During our workshop in May 2008, our participating Advisory Board members identified a critical near-term objective that had not been envisioned in the original proposal: a complete set of identifiers for all geographic features in the *Barrington Atlas of the Greek and Roman World*. The key element here was the urgency of establishing a comprehensive set of stable identifiers, compatible with Pleiades, for historical geographic features so that other projects could begin to incorporate them into their datasets. These IDs could then provide the basis for interaction with Pleiades as its content expands (necessarily more slowly). A scheme for generating these IDs for all *Barrington Atlas* features was devised during the meeting, as was a scheme for a “BATlas ID Resolver,” a web-based service that would accept requests referencing BATlas IDs from outside systems and return links to corresponding Pleiades content items when available. Project team and advisory board members agreed that implementation of these aspects were of the highest priority to ensure broad collaboration beyond the end of the project period. We subsequently

---

<sup>19</sup> <http://www.atlantides.org/trac/pleiades/wiki/PleiadesContent>

implemented both aspects: the *Barrington Atlas* Identifiers themselves, and the supporting "ID Resolver" service. Information about the identifiers was disseminated in a series of blog posts.<sup>20</sup> The resolver works by providing an appropriate explanatory response to any browser that requests a Barrington Atlas ID, as well as a link to corresponding content in Pleiades when available.<sup>21</sup>

An added enhancement, which we had not originally planned, is that the schema for publishing the texts of IRT, and their metadata, enabled the material to be transferred directly into the Epigraphische Datenbank Heidelberg. This is an internationally used resource for finding Latin inscriptions from everywhere outside Italy (which is covered by a separate, parallel, database). This should help scholars to exploit the Libyan material, as well as establishing a new way for other projects to contribute to the database.

### 3 Evaluation, Continuation and Impact

The Concordia project proved rewarding. It is now clear that – given the number of dependencies on delivery timelines for software from external projects and the challenges inherent in software prototyping – that our desire for a working demonstrator at the end of the project were overly ambitious. Nonetheless, tangible results were achieved. We do aim to conduct the demonstration effort later in 2010 or early in 2011, once the PN is fully upgraded and resources permit.

The most important outcome of Concordia is consensus on terms for a thesaurus and conventions for their expression in Atom feeds were agreed between all project participants. This consensus has already resulted in the production of feeds for several of the participating projects and will appear alongside more of them in the near future. The value of typed links (whether expressed in Concordia fashion, open linked data or semantic web contexts) has been highlighted for the ancient studies community; discussions of same are now much more common than they were at the beginning of the project.

Both ongoing and new initiatives are implementing the Concordia approach. ISAW has adopted the use of the vocabulary and of Atom feeds as a common feature across its portfolio of emerging digital publications, including the recently launched Ancient World Image Bank.<sup>22</sup> Pleiades, now supported by new funding from an NEH Preservation and Access Humanities Collections and Reference Resources grant (2010-2013), will be exploiting feeds provided by other partners to effect cross-project discovery from Pleiades place resources. The production of Concordia-conformant feeds is now a standard component of epigraphic publications produced at King's College London and will be a stock component of the Integrating Digital Papyrology software framework. One conclusion of the separately funded (NEH/DFG) Epigraphic Interoperability workshops was the need

---

<sup>20</sup> <http://horothesia.blogspot.com/2008/07/barrington-atlas-ids.html>

<sup>21</sup> An example of the BATlas ID resolver service may be examined here: <http://atlantides.org/batlas/aphroditopolis-aphrodito-77-e3>

<sup>22</sup> Ancient World Image Bank: <http://www.nyu.edu/isaw/publications/awib/>

to equip the Epigraphische Datenbank Heidelberg with Concordia-conformant Atom feeds. A proposal now in preparation for a collaboration between the Alexandria Archive Institute's Open Context system and the American Numismatic Society's nomisma.org initiative builds on the Concordia vocabulary and feed model to facilitate third-party annotation of content in multiple distributed collections.<sup>23</sup>

## 4 Grant Products

- Concordia Project website: <http://concordia.atlantides.org>
- Concordia Atom feeds (specifications and examples):  
<http://www.atlantides.org/trac/concordia/wiki/ConcordiaAtomFeeds>
- Concordia Thesaurus:  
<http://www.atlantides.org/trac/concordia/wiki/ConcordiaThesaurus>
- White paper: Gillies, Sean, *Concordia, Vocabularies, and CIDOC CRM* (New York: Institute for the Study of the Ancient World, New York University, 8 December 2008) <<http://concordia.atlantides.org/docs/concordia-crm.html>>.
- EpiDoc stylesheets for producing Atom feeds from EpiDoc files:  
<http://www.atlantides.org/trac/concordia/wiki/EpiDoc2Atom>

---

<sup>23</sup> Open Context: <http://opencontext.org/>

## 5 Appendices

### 5.1 Project Team

- Professor Roger Bagnall, Director, Institute for the Study of the Ancient World (ISAW), New York University (US Principal Investigator)
- Professor Charlotte Roueché, Department of Classics, Kings College London (UK Principal Investigator)
- Dr. Tom Elliott, Associate Director for Digital Programs, ISAW (US Technical Lead)
- Dr. Gabriel Bodard, Research Associate, Centre for Computing in the Humanities, Kings College London (UK Technical Lead)
- Sean Gillies, ISAW (Chief Engineer)

### 5.2 Advisory Board

Positions indicated were at time of award.

- Rodney Ast (Advanced Papyrological Information System, Columbia University)
- Hugh Cayless (Digital Library Research and Development, University of North Carolina at Chapel Hill)
- Elaine Matthews (Lexicon of Greek Personal Names, Oxford)
- Sebastian Heath (American Numismatic Society)
- Daniel Pett (Portable Antiquities Scheme, British Museum, London)
- Joshua Sosin (Duke Databank of Documentary Papyri, Duke University, Durham, NC)
- Richard Talbert (Ancient World Mapping Center, University of North Carolina at Chapel Hill)